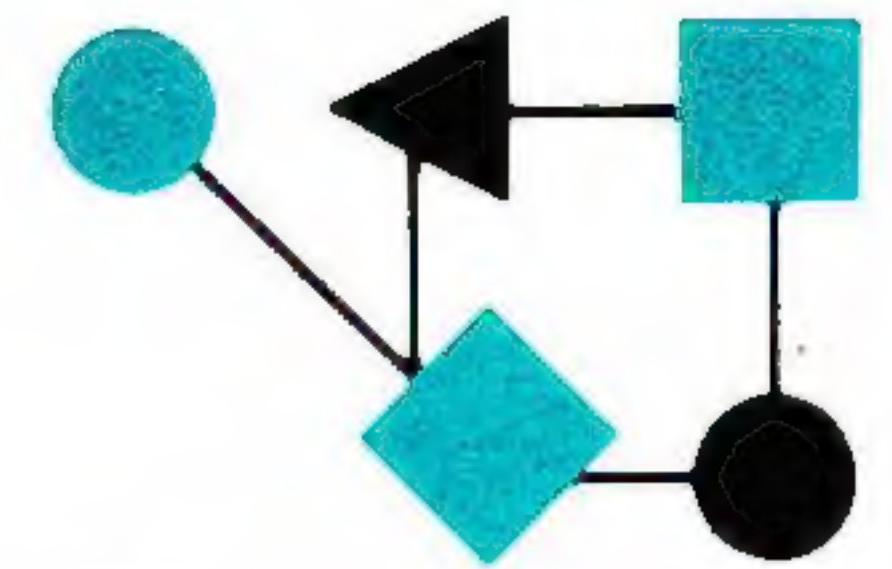


# CONNE~~X~~IONS



## The Interoperability Report

April 1996

Volume 10, No. 4

*ConneXions—  
The Interoperability Report  
tracks current and emerging  
standards and technologies  
within the computer and  
communications industry.*

### In this issue:

Retiring the NSFNET.....	2
SNA.....	12
Announcements.....	21
Book Review.....	26

*ConneXions* is published monthly by Interop Company, a division of SOFTBANK Exposition and Conference Company, 303 Vintage Park Drive, Foster City, California, 94404-1138, USA.

Phone: +1 (415) 578-6900

Fax: +1 (415) 525-0194

E-mail: [connexions@interop.com](mailto:connexions@interop.com)

Subscription hotline: 1-800-575-5717  
or +1 610-892-1959

Copyright © 1996 by Interop Company.  
Quotation with attribution encouraged.

*ConneXions—The Interoperability Report*  
and the *ConneXions* logo are registered  
trademarks of Interop Company.

ISSN 0894-5926

### From the Editor

In addition to working on *ConneXions*, I am involved in the planning and execution of the NetWorld+Interop conference program. One aspect of this work includes planning the *Birds of a Feather* (BOF) sessions. The BOFs are informal after-hour events where participants discuss topics of mutual interest. Judging by recent attendance figures, the most popular topics are Internet Firewalls and Network Management. We have published several articles on Firewalls in recent issues, and will cover this "hot topic" again soon. Meanwhile, we are happy to announce that our May edition will be a Special Issue focusing on Network Management. (For more information about the BOFs just send me e-mail.)

I am also responsible for organizing the *Conference Assessment Team* (CAT) at each of our US NetWorld+Interop events. The CAT program is intended for students who are interested in computer networking. CATs receive complimentary admission to the conference in return for submitting session evaluations. If you know anyone who would be interested in participating, please have them contact me via e-mail. (I prefer the use of e-mail as I am often on the road and unable to respond to phone calls. But with an Internet account it is becoming easier to access e-mail anywhere, be it from a laptop computer, from a "Cyber Café," or from a networked workstation in the office of whomever I am visiting).

Back to this month's issue. Our first article is a look at how the NSFNET was retired from service last spring. For many years, this network formed the core of the Internet and was responsible for many important technological advancements. As previously reported in this publication, the "new" Internet operates as a mesh of service providers interconnected at Network Access Points (NAPs). The article is by Susan R. Harris and Elise Gerich of Merit Network, Inc.

In our *Back to Basics* series we take a look at IBM's Systems Network Architecture (SNA). Many large corporate networks are still based on SNA, but companies are looking for links between their SNA and newer LAN systems. The challenge is to carry SNA traffic on the LAN, and LAN traffic on the SNA network. The article is by Ed Tittel and is adapted from his *PC Networking Handbook*.

We are always looking for new material to publish in *ConneXions*. You can help us by sending your topic suggestions or actual articles via e-mail to: [connexions@interop.com](mailto:connexions@interop.com). Guidelines are available and authors will receive a free lifetime subscription.

## Retiring the NSFNET Backbone Service

*Chronicling the End of an Era*

**by Susan R. Harris and Elise Gerich, Merit Network, Inc.**

### **Introduction**

April 30, 1996, marks the one-year anniversary of the final dismantling of the venerable NSFNET Backbone Service. After more than a year of planning, reconfiguration, shutdowns, and transitions, the U.S. Internet had completed its move to a new architecture composed of multiple backbones, linked at the new interexchange points.

The midnight NSFNET shutdown went remarkably smoothly, as did most of the events leading up to the final phaseout. This article looks back on the timelines, dependencies, delays, emergencies, and successes that marked the final year of the NSFNET. We begin by taking a brief look at the history of what was the world's largest and fastest network for research and education.

### **A Brief History of the NSFNET**

The *National Science Foundation* inherited the responsibility for nurturing the U.S. Internet from the *Advanced Research Projects Agency* (ARPA). From its inception in 1985–1986, the NSFNET program laid the foundation of the U.S. Internet and was the main catalyst for the explosion in computer networking around the world that followed. [1] The first NSFNET, a 56Kbps backbone based on LSI-11 Fuzzball routers, went into production in 1985 and linked the six nationally funded supercomputer centers (the five NSF centers and the National Center for Atmospheric Research). Soon after the network's inception, the need for more advanced networking technology was indicated when rapid growth in traffic precipitated serious network congestion. In 1987, NSF issued a competitive solicitation for provision of a new, faster network service. The new service would provide a network backbone to link the six supercomputer centers and seven mid-level networks. The mid-level networks would in turn connect campuses and research organizations around the country, creating a three-tiered network architecture that remained in place until the end of the NSFNET backbone service.

In fall 1987, NSF selected Merit Network, Inc., and its partners MCI, IBM, and the State of Michigan to manage and re-engineer the new backbone service. Eight months after the NSF award, the NSFNET partnership delivered a new T1 backbone network that connected 13 sites: Merit, NCAR, BARRNet, MIDnet, Westnet, NorthWestNet, SESQUINET, SURAnet, and the NSF supercomputer centers. Two additional regional networks, NYSERNet and JVNCnet, were also served by the backbone, because each was co-located at a supercomputer center. Each of the 13 backbone nodes, known as *Nodal Switching Subsystems*, was composed of nine IBM RTs linked by two Token Rings with an Ethernet interface to attached networks. There were 14 T1s connecting the sites, on which a virtual topology was constructed. Each virtual path represented one-third T1 to the site.

In 1989 the backbone was re-engineered, increasing the number of T1 circuits so that each site had redundant connections to the NSFNET backbone as well as increasing router capability to full T1 switching. With this upgrade, the NSFNET's physical topology equaled its virtual topology. By then, the traffic load on the backbone had increased to just over 500 million packets per month, representing a 500% increase in only one year. Every seven months, traffic on the backbone doubled, and this exponential growth rate created enormous challenges for the NSFNET team. [1]

## Upgrade to T3

To handle the increase in traffic, Merit and its partners introduced a plan to upgrade the backbone network service to T3. The NSF also wanted to add a number of new backbone nodes, and asked Merit to prepare proposals for the added cost of new nodes at T1 and T3 speeds, while the NSF issued a solicitation to the community for those interested in becoming new NSFNET sites. It was eventually decided by the NSF that the partners would increase the total number of backbone nodes on the NSFNET from 13 to 16, all running at 45 Mbps. Additional sites served by the T3 NSFNET backbone service would include Cambridge MA (NEARNET), Chicago's Argonne National Lab, and Atlanta GA (SURAnet).

In late May 1990, Merit's cooperative agreement with NSF was modified to cover the additional work. By the end of the year, Merit, SDSC, and NCSA were connected to an early T3 service and began testing the new T3 routers with real traffic. In addition, a new T3 research and test network was implemented to parallel the existing T1 test facility.

Important architecture and equipment changes came with the new T3 network. The core backbone equipment was moved from the universities and supercomputer sites to MCI's points-of-presence (POPs), and the RTs were replaced with RS/6000s and a card-to-card forwarding architecture. Many of the techniques introduced in the T3 RS/6000 routers have since been adopted by commercial router vendors.

As the backbone network service was growing in complexity and was re-engineered, increasing focus and resources were needed to keep pace with more complex technical, business, and policy environments. To meet these organizational challenges, ANS was created and announced in September 1990. ANS began to provide service for NSFNET as a subcontractor to Merit, with IBM, MCI, and others continuing to infuse new technology to develop the infrastructure.

During 1991, a year of refining the new backbone technology, the T1 and T3 networks existed in parallel. Difficulties in tuning the new technology prevented the network from being moved to full production status until late in the year, when all sixteen backbone sites comprising the NSFNET service were connected to the new ANSnet national T3 infrastructure. With expansion work completed and improved performance validated, several sites began using the T3 for their primary traffic path by November 1991. A final round of testing in mid-December set the stage for moving the remaining NSFNET traffic to the new backbone service in early 1992. The network now exceeded the T1 structure in stability by a factor of ten, with fewer outages and errors in all categories.

The upgrade of the NSFNET backbone service to T3 was not only a technological and organizational challenge of the highest order. It also precipitated a greatly-needed, though contentious, community dialogue about the evolution and commercialization of the U.S. Internet. *Internet Service Providers* (ISPs) were springing up all over the country, from local dial-up providers to larger companies providing T1 and eventually T3 service, and there were now a growing number of vendors offering TCP/IP networking products and services.

During 1992, the National Science Board authorized an extension of Merit's cooperative agreement for eighteen months beyond the October 1992 expiration date in order for NSF to develop a follow-on solicitation for national networking.

## Retiring the NSFNET Backbone (*continued*)

This solicitation was one that would accommodate the growing role of commercial providers and allow NSF to step back from actually operating a network to concentrate on supporting leading-edge research initiatives. NSF published a draft solicitation for community comment in 1992, and a new solicitation was issued in May 1993.

Early in 1994, awards for building the new architecture were given to Merit and USC's Information Science Institute for the *Routing Arbiter* service, to MCI for the *very high speed Backbone Network Service* (vBNS), and to three providers for the *Network Access Points* (NAPs): Sprint, MFS Datanet, and Bellcore, representing Ameritech and Pac\*Bell. NSF also awarded Merit a transition extension that began in May 1994 and lasted until April 1995, when the NSFNET backbone service would be retired and all connections would be switched to a new service.

### Deadlines and Commitments

Moving the U.S. Internet to a new architecture in the months between the 1994 awards and the April 30, 1995 termination date was a frightening challenge for the regional networks, the ISPs, and the NSFNET partnership. Before the backbone could be decommissioned, four main tasks had to be accomplished by the networking community:

- Establish the NAPs and move them to production status.
- Attach to the NAPs the NSFNET and the ISPs that provided service to the regionals.
- Develop the RA Service by placing Route Servers at the NAPs and setting up a routing registry.
- Move the regionals off the NSFNET and attach them to networks operated by ISPs.

According to NSF's ambitious transition schedule, the new NAPs would be available by August 15, 1994. The NSFNET backbone service would then attach to the NAPs, with all current attachments to the NSFNET remaining in place. The ISPs would then begin to attach to the NAPs, and regional networks that attached to NSFNET would begin to establish connections to the ISPs. By October 31, the regionals would cut over all traffic to the ISPs and disconnect their attachments to the NSFNET. Only the supercomputer centers would remain attached to the NSFNET. The vBNS would be deployed by January 1, 1995, and attached to the NAPs by February 1, 1995.

As it turned out, all of these actions were delayed, and revised deadlines established.

### Establishing the NAPs

The first Network Access Point to go into production was the Washington, D.C. NAP (MAE-East, the Metropolitan Area Ethernet). MFS had been operating MAE-East since 1992, and MAE-East had served as a model for the NAPs as defined in NSF's solicitation. In fall 1994, MAE-East was upgraded from a 10Mbps Ethernet to FDDI; internetMCI and SprintLink, which had already attached to the MFS facility, upgraded their connections to FDDI, as did the NSFNET.

The Sprint NAP, a bridged FDDI/Ethernet hybrid, was up and running by the end of the summer; ANSnet/NSFNET, SprintLink, and internetMCI attached to it in September.

The Sprint and Washington, D.C. NAPs began to carry much of the traffic for the U.S. Internet once networks began to move off NSFNET in November 1994, because the Pac\*Bell and Ameritech ATM NAPs were still being deployed and went into production several months later. Both facilities were physically in place by October 1994, but problems with ADSU performance and a concern with ATM switch buffer sizes led to a lack of confidence in the ability of the ATM NAPs to sustain the traffic load.

As a result, both Pac\*Bell and Ameritech decided to deploy interim configurations, and put FDDI LANs into production in March 1995. Some ISP routers on the FDDIs at these contingency NAPs were also connected to DS3 ports on the ATM switch, so they could pass traffic across the FDDI while still transmitting to ATM-connected peers. As of January 1996, this infrastructure is still in place at the Pac\*Bell and Ameritech NAPs.

## Deploying the Route Servers

The Routing Arbiter service has two main components: the *Route Servers*, SPARC 20s deployed at the NAPs, and the *Routing Arbiter Database*, successor to the Policy Routing Database used to configure the NSFNET backbone service. [3, 4, 5, 6, 7]

In November 1994, primary and backup Route Servers were shipped from ISI to each of the NAPs. Once the necessary data circuits, front-end systems, controllers, ATM switches, and FDDI bridges were installed and tested, addressing schemes worked out, security procedures implemented, and 24/7 network monitoring in place, the Routing Arbiter team began to set up peering sessions with customer routers at the NAPs. Out-of-band access—a prerequisite for declaring the Route Servers fully in production—became available several months later.

By April 1995, the Route Servers were peering with more than a dozen providers at the Sprint and Washington, D.C. NAPs. In July, production RA services were announced at the Sprint NAP, and announcements for the other NAPs soon followed. At each exchange point, the Route Servers began importing and exporting routes to numerous ISPs. The ISPs maintained sessions with other peers as well as the Route Servers, comparing the routing information from both sessions for consistency.

## NACRs and the PRDB: The Long Goodbye

Merit originally planned a December 1994 retirement for the Policy Routing Database (PRDB), which had been used to configure the NSFNET's backbone routers since 1989. The PRDB would be replaced by the Routing Arbiter Database, which would then become part of the Internet Routing Registry (IRR) along with the RIPE NCC, MCI, ANS, and CA\*net registries. The IRR would be an important global resource—a public repository of announced routes and routing policy in a common format, so that ISPs could use the information stored in any and all registries to configure their backbone routers, analyze routing policy, and build tools to help in these efforts.

The PRDB was established to maintain information about what were considered legitimate destination announcements from the various regionals. The primary goal of maintaining this information was to prevent routing loops. When BGP replaced EGP as the inter-domain routing protocol in 1994, suppression of routing loops no longer had to be so administratively controlled. The information in the PRDB was then mainly used to record routing policies such as path preferences and to generate the backbone configuration files.

## Retiring the NSFNET Backbone (*continued*)

NSF's follow-on solicitation for the new architecture specified a continuation of the function that the PRDB played in the T1/T3 NSFNET. The goal was to record global routing policy information based on each Autonomous System's policy. RIPE had pioneered this work in the European arena, and the data exchange format described in RIPE-181 (RFC 1786) was adopted as the "standard" for Internet Routing Registries. [2] The RADB adheres to this model.

The challenge was to establish and populate the RADB before the retirement of the NSFNET and the PRDB. By summer 1994, the RIPE NCC registry had been in production for two years, and CA\*net and internetMCI were creating routing registries to support their customers. ANSnet would continue to use the PRDB until the RADB was established. But the dilemma was how to transition from NSFNET-centric information to the AS-specific information needed for the RADB, while continuing to provide a stable router configuration environment for the NSFNET service.

Merit's December target date for retiring the PRDB was based on the assumption that the regionals would be off the backbone by October 31. When it became clear that they weren't going to make that deadline and the PRDB would need to continue to support the NSFNET and its regionals well beyond the end of October, a plan was proposed to transition to the RADB to support the NSFNET in its last months.

The new situation presented several problems. First, the tools used to configure the NSFNET/ANSnet routers were based on PRDB attributes, not RIPE-181. Second, the RADB was not yet populated with data. And finally, the PRDB described AS690 policy with respect to its peer ASs on a per-prefix basis; in the RIPE syntax, the basis for describing routing policy was the Autonomous System where the route originated. With more than 40,000 prefix-based policies for the regionals, the PRDB was used to generate about 100 configuration files of around 250,000 total lines every two weeks, and those policies needed to be re-expressed in a RIPE-compatible format.

Continuing the Policy Routing Database for long-term support of ANSnet was inadvisable. If ANS continued to use the PRDB for AS690 routing after the transition, the PRDB's non-standard format would create a barrier to sharing global routing policies and building tools to aid with global routing. A solution had to be found that would provide stable routing through the transition, and, once the NSFNET was retired, allow the ANS registry to take its place alongside the other registries in the IRR.

To solve the problem, Merit proposed a modification to RIPE-181—a temporary attribute that would specify the peer or adjacent AS announcing the route to AS690. The community agreed to Merit's proposal, and the new expression came to be known as the "advisory attribute." Merit now needed to quickly modify the PRDB configuration tools so they would generate the new attribute, populate the RADB with the data needed to generate AS690 configuration files, and make sure that the new configurations exactly matched those produced by the PRDB.

By December 1994, all the data in the PRDB had been converted to RIPE-181-style expressions and entered in the RADB. By February, the RADB had been populated with RIPE-181-style Maintainer and AS Objects. The databases were running in parallel, with changes to the PRDB automatically reflected in the RADB.

Other organizations whose routing information wasn't related to the NSFNET were also populating the RADB throughout the winter of 1994–95; this was another variable that had to be accommodated as the new database emerged.

Finally came the painstaking task of comparing the config files generated by each database. Merit's Dale Johnson went over the large, quarter-megabyte files line by line, adjusted the configuration tools to compensate for any differences in net lists, and repeated the process over and over until the configs matched perfectly. The RADB finally replaced the PRDB a week after the NSFNET was retired.

### Moving the Regionals off NSFNET

NSF and Merit coordinated the process of moving the regionals to new Internet Service Providers, with Merit taking the lead in planning the transition. NSF's new Inter-Regional Connectivity program helped support new attachments not only for NSFNET peer networks—regionals like SURAnet and NYSERNet that connected directly to the NSFNET backbone—but also to downstream networks such as NevadaNet and MOREnet. Most of the regionals selected internetMCI or SprintLink as their ISP; CERFnet set up its own ATM connection to each NAP.

In mid-October 1994, NSFNET Program Director Priscilla Huston sent a letter to the regionals asking them to send a transition calendar and engineering overview to Elise Gerich of Merit and to her. Huston also asked the regionals to notify her if they weren't going to make the October 31 deadline for moving off the NSFNET.

As it turned out, none of the networks made it. MOREnet, one of the downstream regionals, missed by only a day; other networks slipped by as much as three or four months. The first NSFNET peer network to make the transition was CA\*net, which faced a hard deadline from its link provider for terminating its connection to the NSFNET. The other cutovers were pushed back because of delays in provisioning the ISPs selected by the regionals, and because of reticence on the part of the regionals to move off the NSFNET backbone service.

On one or two occasions, networks that had made the transition had to pull back to full NSFNET connectivity because of deployment problems on the new ISP backbone. In general, though, once the regionals had selected an ISP and completed all the testing, re-routing, and re-configurations necessary to make the switch, traffic flowed smoothly over the new infrastructure.

### 60-Day Notices: No Turning Back

Early in January, when SURAnet notified NSF and Merit that it was ready to move off the backbone, Merit sent ANS the first message to dismantle NSFNET backbone service—a 60-day termination notice for ENSS 138 in Atlanta. The ENSSs (*Exterior Nodal Switching Subsystems*) were installed at regional networks attached to the NSFNET, and acted as end nodes for the backbone. This and subsequent termination notices were irrevocable; once sent, there would be no more NSFNET service through that node.

Later in January, NYSERNet and the Cornell Theory Center notified NSF and Merit that they were ready to terminate their NSFNET attachments. The other regionals and supercomputer centers followed suit, one by one, as the April deadline neared.

## Retiring the NSFNET Backbone (*continued*)

On February 28, Gerich sent Jordan Becker of ANS the formal, 60-day notice for termination of the NSFNET backbone service at 19 locations:

ENSS 128	Palo Alto	April 30, 1995	midnight PST
ENSS 129	Champaign	April 30, 1995	midnight CST
ENSS 130	Argonne	April 30, 1995	midnight CST
ENSS 131	Ann Arbor	April 30, 1995	midnight EST
ENSS 132	Pittsburgh	April 30, 1995	midnight EST
ENSS 133	Ithaca	April 30, 1995	midnight EST
ENSS 134	Cambridge	April 30, 1995	midnight EST
ENSS 135	San Diego	April 30, 1995	midnight PST
ENSS 136	College Park	April 30, 1995	midnight EST
ENSS 137	Princeton	April 30, 1995	midnight EST
ENSS 139	Houston	April 30, 1995	midnight CST
ENSS 140	Lincoln	April 30, 1995	midnight CST
ENSS 141	Boulder	April 30, 1995	midnight MST
ENSS 142	Salt Lake City	April 30, 1995	midnight MST
ENSS 143	Seattle	April 30, 1995	midnight PST
ENSS 144	Moffett Field	April 30, 1995	midnight PST
ENSS 145	College Park	April 30, 1995	midnight EST
ENSS 146	DC	April 30, 1995	midnight EST
ENSS 147	MFS	April 30, 1995	midnight EST

The list included the NSFNET attachments at the NAPs, which were coexistent with ANSnet. ENSS 138 in Atlanta wasn't included, since a termination notice for that node had been issued earlier.

By March, backbone traffic had declined dramatically, but not quite as fast as NSF and Merit had expected. Gerich, concerned that the regionals and ISPs weren't moving fast enough, sent e-mail to the community noting that a significant amount of traffic was still traversing the NSFNET. Merit posted a histogram showing the top 10 originators of traffic into the backbone in February 1995, and reminded networks attached to nodes highlighted on the graph about the April 30 deadline.

Later that month, Merit discontinued the T1 safety net that had backed up the T3 infrastructure since 1992.

By the middle of April, only seven regionals had completely severed their ties to the NSFNET backbone service. Other networks had cut over to a new service provider, but continued to peer with the NSFNET for backup purposes. As the final deadline neared, Merit and the NSFNET Executive Committee became concerned that these redundant connections would make it difficult to identify outstanding reachability issues before the April 30 cutoff.

To spot any pockets of unreachable destinations before it was too late, Merit on behalf of the NSFNET Executive Committee notified the NSFNET community on April 14 that it would terminate peering sessions with all organizations still attached to the NSFNET Backbone service at 9:00 a.m. on Friday, April 21. On April 28, all sessions with the NSFNET service would be permanently terminated; ANS would terminate operation of the NSFNET Backbone service on April 30.

This announcement created quite a stir among the networks attached to the backbone. Several said that they'd lose their Internet connectivity completely if their NSFNET peering was shut down before April 30, and requested that their session stay up.

One provider was still relying on his MAE-East NSFNET connection for all his East coast traffic; another requested clemency for a non-production peer router that was proving essential for network diagnostics. A Midwest network's installation of a T3 circuit had been delayed; the operators weren't concerned about reachability if their NSFNET peering was shut down, but about capacity—a large volume of traffic was still traversing the NSFNET, and cutting back to a T1 would lead to unacceptable response times. Merit made separate arrangements to accommodate each network, but held to the new deadline.

As it turned out, the test shutdown had to be postponed. In the early morning hours of April 21, Merit notified the community that it would have to delay the regular Friday backbone configuration run. The volume of routing configuration changes had increased so dramatically as networks switched to new providers that some of the files grew large enough to truncate during production, and produced corrupt configuration files. Merit wasn't confident that it would be able to produce complete and correct configurations in time for the normal 8:00 a.m. configuration window. Additional file space had to be allocated before the configs could be run, and Merit needed to work with several ASs to reduce the number of net lists in the config file. This meant postponing the Friday shutdown until Saturday, and delaying the NSFNET discontinuation until Tuesday, April 25. The test shutdown had indeed pinpointed at least one problem as a result of delayed transitions: the processing of several thousand simultaneous changes to router configurations was more than the PRDB could handle.

On Monday, one network jumped the gun, and surprised Merit and ANS by turning off its ENSS. The ANS staff noted that no harm had been done, but reminded the sysadmin that the plan was to manually turn off peering on the 25th, and shut off the ENSSs on the 30th. IBM was to physically remove the routers beginning in May.

On April 25, the peering sessions on 15 ENSSs were commented out of the configuration files and the NSFNET Backbone Service was, for all intents and purposes, terminated. The next Sunday evening at midnight, a dozen or so staff from Merit and ANS gathered in the University of Michigan NOC to turn off the ENSSs, one by one, at midnight in each respective time zone. One or two regional operations centers called the ANS NOC about unreachable ENSSs, but "mostly the NSFNET went away silently," as one ANS engineer remarked, "or rather, with only the sound of drives and fans spinning down in distant machine rooms."

On May 8, with Merit confident that the RADB was producing consistent configuration files for the ANSnet and ANS ready to take over configuration generation for AS690, the PRDB made a graceful exit. The new architecture was in place: internetMCI and SprintLink had absorbed the NSFNET regionals as their customers; the RADB and the databases maintained by the RIPE NCC, internetMCI, CA\*net, and ANSnet had replaced the PRDB as a means of describing routing policy.

Farewell NSFNET! And congratulations to the hundreds of people who helped make the backbone such a great success.

## Retiring the NSFNET Backbone (*continued*)

### References

- [1] Frazer, Karen D., "NSFNET: A Partnership for High-Speed Networking," Merit Network, Inc., 1995. Available from URL: <http://www.merit.edu/nsfnet/final.report>
- [2] Bates, T., Gerich, E. Joncheray, L., Jouanigot, J-M., Karrenberg, D. Terpstra, M. and Yu, J., "Representation of IP Routing Policies in a Routing Registry (ripe-81++)," October 1994. Subsequently published as RFC 1786, March 1995.
- [3] D. Estrin, J. Postel, and Y. Rekhter, "Routing Arbiter Architecture," *ConneXions*, Volume 8, No. 8, August 1994.
- [4] Yu, J., Chen, E. and Joncheray, L., "A Routing Design for the Initial ATM NAP Architecture," *ConneXions*, Volume 8, No. 11, November 1994.
- [5] Yu, J., Chen, E. and Joncheray, L., "Routing for the Initial ATM-NAPs In the Presence of the Special ANS/NSFNET Attachment," *ConneXions*, Volume 9, No. 1, January 1995.
- [6] Yu, J., "The RA Route Server Service Overview," *ConneXions*, Volume 9, No. 8, August 1995.
- [7] Manning, B., "The Routing Arbiter in the post-NSFNET Service World," *ConneXions*, Volume 9, No. 9, September 1995.
- [8] CISE/NCR, "NSF 93-52 — Network Access Point Manager, Routing Arbiter, Regional Network Providers, and Very High Speed Backbone Network Services Provider for the NSFNET and the NREN™ Program," Program Guideline nsf9352, May 1993.
- [9] J. Scudder and S. Hares, "NSFNET Traffic Projections for NAP Transition," NANOG Presentation, October 1994, URL: <http://www.merit.edu/routing.arbiter/NANOG/Scudder-Hares.html>
- [10] RIPE NCC, "Delegated Internet Registry in Europe," RIPE-NCC F-2 Version 0.3, March 1994.
- [11] S. Williamson and M. Kosters, "Referral Whois Protocol," RFC 1714, November 1994.
- [12] C. Alaettinoglu, and J. Yu, "Autonomous System Path Expression Extension to RIPE-181," Technical Report, USC/ISI, March 1995.
- [13] T. Bates, M. Terpstra, D. Karrenberg, "The PRIDE project directory," March 1994, URL: <ftp://ftp.ripe.net/pride/>
- [14] Govindan, R., Alaettinoglu, C., Varadhan, K., and Estrin, D., "A Route Server Architecture for Inter-Domain Routing," USC Technical Report # 95-603, January 1995.
- [15] Merit Network Inc., "Routing Arbiter for the NSFNET and the NREN, First Annual Report," April 1995.
- [16] *ConneXions*, Volume 3, No. 8, August 1989, "Special Issue: Internet Routing."
- [17] *ConneXions*, Volume 5, No. 1, January 1991, "Special Issue: Inter-domain Routing."

- [18] Radia Perlman, *Interconnections: Bridges and Routers*, Addison-Wesley, Reading, Massachusetts, 1992.
- [20] Estrin, D., Rekhter, Y., Hotz, S., "A Unified Approach to Inter-Domain Routing," RFC 1322, May 1992.
- [21] Estrin, D., Rekhter, Y., Hotz, S., "Scalable Inter-Domain Routing Architecture," SIGCOMM 1992, Volume 92, No. 4, October 1992.
- [22] Estrin, D., Li, T., Rekhter, Y., "Source Demand Routing Protocol: A Component of The Unified Approach to Inter-Domain Routing," *ConneXions*, Volume 6, No. 11, November 1992.
- [24] Karrenberg, Daniel, "The RIPE NCC and the Routing Registry for Europe," *ConneXions*, Volume 7, No. 11, November 1993.
- [25] Stockman, B., "Global Connectivity: The Global Internet Exchange (GIX)," *ConneXions*, Volume 7, No. 11, November 1993.
- [26] Rekhter, Y., and Li, T. (Editors), "A Border Gateway Protocol 4 (BGP-4)," RFC 1654, July 1994.
- [27] NNSC Staff, "Profile: NSFNET," *ConneXions*, Volume 1, No. 2, June 1987.
- [28] Hans-Werner Braun, "The new NSFNET backbone network," *ConneXions*, Volume 2, No. 12, December 1988.
- [29] Laura Kelleher, "Merit/NSFNET Information Services," *ConneXions*, Volume 3, No. 6, June 1989.
- [30] P. Bosco & H.-W. Braun, "The NSFNET T1/T3 Network Celebrates 5 Years," *ConneXions*, Volume 7, No. 9, September 1993.
- [31] Bill Norton, "Network Discovery Algorithms for the NSFNET," *ConneXions*, Volume 8, No. 4, April 1994.
- [32] Dern, D., "The ARPANET is Twenty: What We Have Learned and The Fun We Had," *ConneXions*, Volume 3, No. 10, October 1989.

**SUSAN R. HARRIS** is a technical writer in the Merit Internet Engineering group. Before she came to Merit in 1992, she edited the University of Michigan's campus computing newsletter for eight years. Harris holds an M.A. and Ph.D. in Ancient Near Eastern History from the University of Michigan. Her networking publications include "Whois Tool Provides Wealth of Information," *OnTheInternet*, May/June 1995. E-mail: [srh@merit.edu](mailto:srh@merit.edu)

**ELISE GERICH** is Merit Associate Director for National Networking and a co-Principal Investigator for the Routing Arbiter project. She came to Merit in 1987 as a member of the original NSFNET team. Gerich's responsibilities for the NSFNET focused on technical interfaces and policy interactions between mid-level networks, sister Federal agencies, and international network research organizations. She serves as North American co-chair of the Internet Engineering and Planning Group and is a member of the IAB. Publications include "Guidelines for Management of IP Address Space," RFC 1466; "Unique Addresses are Good," RFC 1814; "Management and operation of the NSFNET backbone," *Computer Networks and ISDN Systems*, Volume 23 (1991); and see reference [2] above. E-mail: [epg@merit.edu](mailto:epg@merit.edu)

## Back to Basics: SNA

by Ed Tittel

### Introduction

A lot of LAN managers know nothing about IBM's *Systems Network Architecture* (SNA). There's a good reason for this. SNA started out in the mainframe world, and mainframe gurus haven't been very forthcoming about sharing their information. SNA networks are essentially networks of mainframe and mini-mainframe computers and devices.

Like most other mainframe technologies, SNA networks have traditionally been managed and run from a centralized location. Lots of companies that use LANs made up of PCs still have SNA networks up and running. Nevertheless, there's no communication, except a very complicated "sneakernet" process, between the two systems. More and more companies are looking for links between their SNA and LAN systems. The network administrator has to find a way to carry SNA traffic on the LAN, and LAN traffic on the SNA network.

Migration to LANs and WANs often means moving from legacy mainframe systems to distributed network-based architectures, but that invariably involves a long period of transition. One of the main challenges that faces a LAN or WAN administrator is connecting an SNA network with disparate LAN systems. In this article, we'll try to point out important terminology, equipment, and techniques for bringing together the LAN and the mainframe networking worlds.

By the time you finish this article, you should have a basic understanding of SNA. We hope you'll also have more insights for solving typical LAN-to-host connection problems. The scope of this article can't include every detail about interconnecting SNA networks and LAN systems, but we invite you to consider it a good starting point. While you're considering, be advised that if your organization does not have any IBM mainframes or minicomputers, you can consider yourself lucky!

### SNA Architecture

The SNA architecture is a host-centric network system arranged in a hierarchical topology. Dumb terminals and other devices communicate with the host using private networks and packet-switched services. The host maintains the brains and the processing power, while the dumb terminals serve as screen displays for its users. Typically, mainframe host management and control is centralized because the host is itself in one place, but the centralized management model also offers a cost-effective way to deal with expensive resources.

Today, that structure is disintegrating. The evolution of LANs and the deployment of client-server-based networks has caused an erosion of central control. In client-server networks, the client plays an important role in the processing of programs, instead of serving as a dumb terminal. The processing is therefore distributed instead of host centric. Obviously, distributed processing demands new technologies and architectures. The mainframe is nowhere near dead and buried, so the two worlds must also be able to get along.

### SNA History

The *International Organization for Standardization* (ISO) began its development of *Open Systems Interconnection* (OSI) standards 20 years ago and published its seven-layer model in 1979. The OSI model is based on a layered approach to networking just like its predecessor, SNA. While organizations were still waiting for the OSI standards to emerge, IBM's SNA architecture was able to slip into large organizations and gain a strong foothold in the industry.

Because of the high costs of SNA, which involves mainly mainframes and stratospheric prices, most small organizations were not able to afford SNA networks, but many of them rented access to time-sharing networks that permitted them to partake of SNA services and applications without absorbing all of the associated costs. Even so, although “big iron” isn’t always a “big company” phenomenon, big iron expertise is much more common in larger organizations than in smaller ones.

### SNA Releases

SNA was designed to connect various devices like terminal controllers, microcomputers, minicomputers, and mainframe computers. SNA began as an architecture in 1974. There have been several subsequent versions of SNA, most recently the *Advanced Peer-to-Peer Networking* (APPN) specification, which is also called the *new SNA*. The following list represents IBM SNA and APPN releases over the years along with a brief description, where the kind of hierarchical network organization that SNA supports is called a *tree structure*:

- 1974: Original release of a tree structure that provided support of only one host and its terminals.
- 1976: A tree structure that allowed for multiple hosts each with its own trees, with communications between trees permissible by communicating through the root of each tree.
- 1979: This provided for more general communications and eliminated the need for communications having to go through the root of each tree.
- 1985: Certain topologies of LANs and hosts were supported.
- 1986: APPN routing was introduced on System/36.
- 1988–1991: Some IBM systems could single-hop to adjacent nodes, AS400 could multihop beyond adjacent nodes.
- 1991–1993: APPN routing capability delivered to AIX and SAA environments.
- 1993: APPN delivered to mainframe via free upgrade of VTAM and NCP.

### Piecing Together the Jigsaw Puzzle

IBM had a near lock on the computer world before the explosion in LAN computing occurred in the 1980s. One consequence of its dominance is the large number of corporate networks that started out as SNA networks. Many companies started out with an IBM SNA system, but added LAN technologies to their enterprise networks over time.

Today, a major puzzle for these organizations is how to integrate their existing and evolving SNA network, and the newer but rapidly expanding LANs. Many companies still maintain two separate WANs, one of them for connecting mainframes and devices using SNA, and the other for connecting LANs using PC technology.

IBM introduced Advanced Peer-to-Peer Networking (APPN) to aid in linking its customers in an SNA network to an enterprise networking environment that incorporated dissimilar networks. APPN has never been the solution it was intended to be. APPN didn’t perform well in internetworking, but provided for much more efficient internetworking between SNA networks. One of APPN’s greatest contributions is that it changes SNA from its hierarchical structure to a peer-to-peer structure.

*continued on next page*

**SNA (*continued*)****IBM Wakes Up!**

In 1994, IBM began announcing products and strategies that would prove it was serious about connecting its own SNA products to the rest of the world. We'll discuss some of these products along with other vendors' solutions later in this article. The industry is still debating whether IBM waited too long to jump on the LAN bandwagon. APPN took a long time to step up to the plate, and some organizations simply couldn't wait. Some say that SNA and APPN are dying, yet others claim that SNA is regaining strength. We're not sure whether it's poised for higher heights or ready to croak, but either way, it's not dead yet, and must be dealt with.

**SNA Briefing**

SNA was designed to operate at the data link layer (layer 2) of the OSI model. It is a connection-oriented protocol that predetermines the path that data will travel through the network before data are sent. The protocol is designed so that the communications devices on the network request a connection or session with the receiving node and then make sure the data are sent. The timing of transmissions involved in SNA messaging makes SNA unsuitable for transmission in a LAN-based environment. We'll talk more about that later in the article. In the meanwhile, here's an introduction to important SNA terminology and concepts.

**Nodes**

SNA has different components in its architecture. These components are defined as *nodes*, which are physical points in an SNA network that contain one or more network components. Any device that conforms to SNA's specifications and houses SNA components can be considered an SNA node. There are four types of SNA nodes:

- *Type 1*: Terminals—devices that interface with a user.
- *Type 2*: Controllers—cluster devices that manage multiple terminals.
- *Type 3*: Ooops, there isn't a type 3 node!
- *Type 4*: Front-end Processors—devices that manage communications with a host and controllers; FEPs poll the controllers periodically.
- *Type 5*: Hosts—the brains of the network (host means an autonomous computer, which in the IBM world is often synonymous with mainframe).

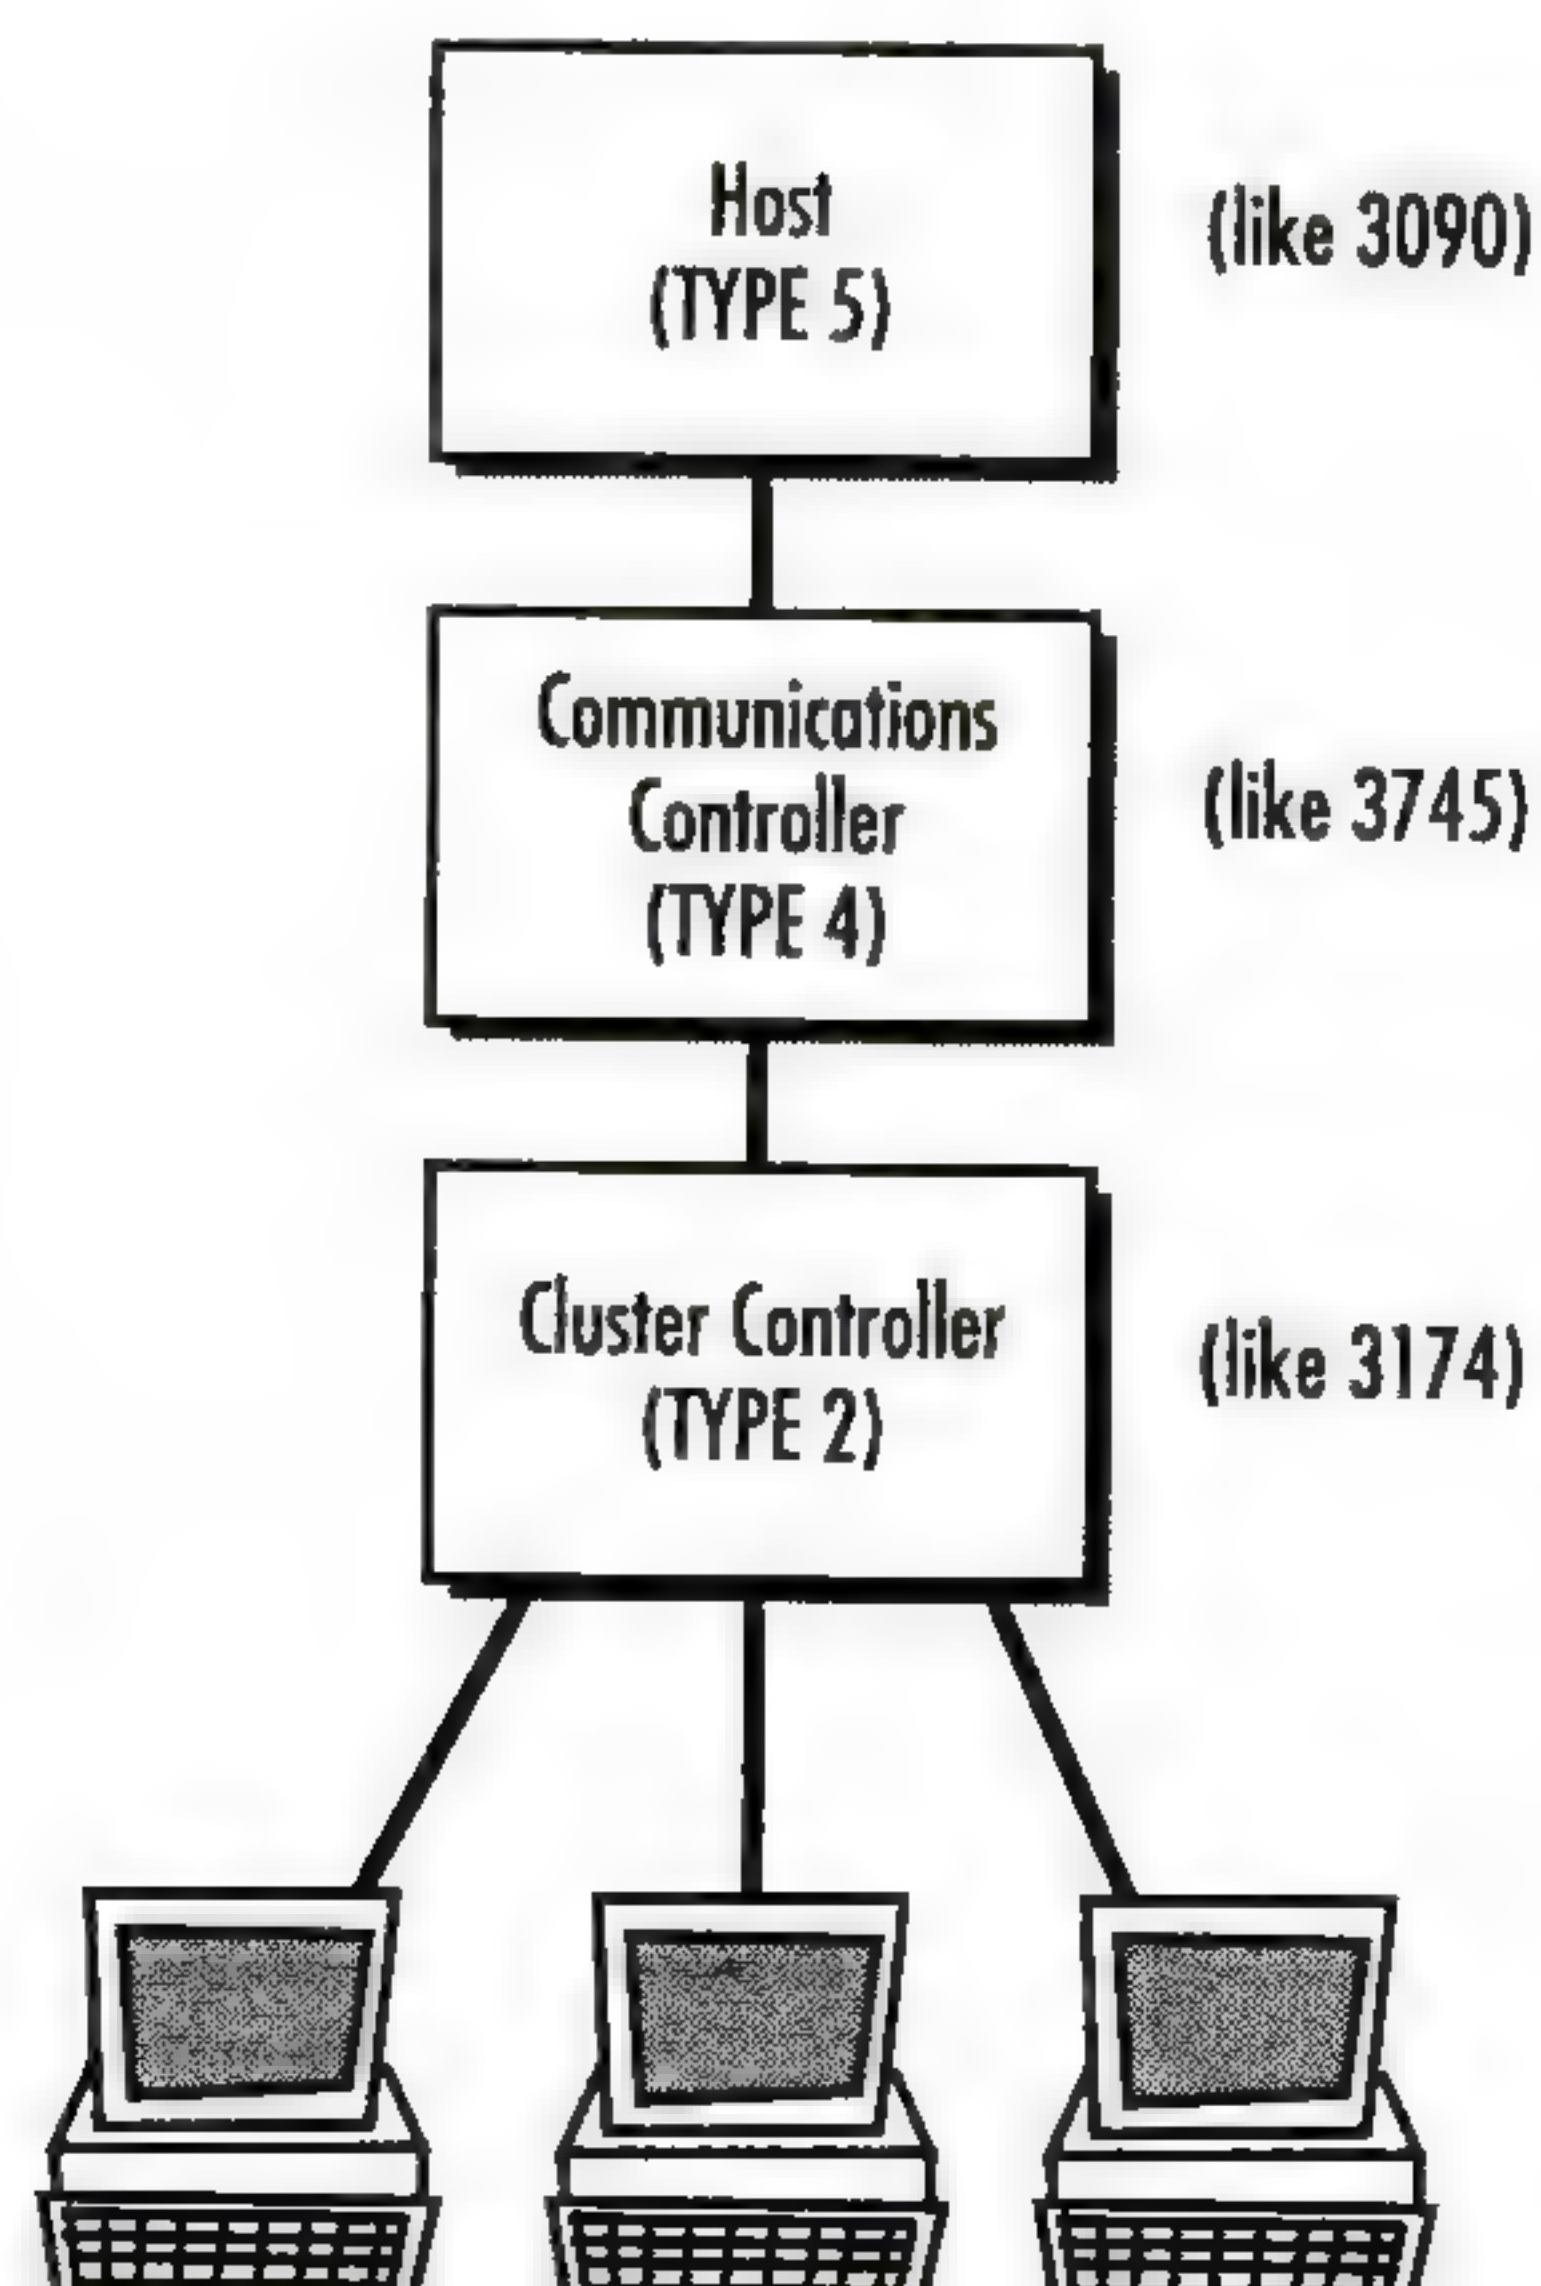


Figure 1: Basic node configuration in an SNA network

Figure 1 depicts a basic node configuration. At first glance, you'll probably notice the tree structure or hierarchical topology for these nodes. This comes from the original design of SNA, which required data to travel up and down this tree structure to go from sender to receiver.

IBM's 3174/3274 controllers were typically connected via dedicated SDLC WAN links (9.6–56Kbps serial connections, often leased telephone lines) to an IBM front-end processor at the mainframe site. Dedicated lines guarantee bandwidth in an SNA network but don't allow commingling of SNA with LAN-based traffic. These dedicated links were important to SNA because of its time-sensitive nature. The FEP's most important job, therefore, was to poll controllers and wait for a response, usually at 10-second intervals. If it did not receive a response from a controller within this time frame, the FEP assumed the controller was down and terminated the session.

### Physical Units

As the name suggests, physical units (PUs) are the physical devices or hardware that can be found in an SNA network. Each SNA node contains at least one PU along with its associated services to manage the physical device. PUs are given the same designations as their corresponding node types:

- PU Type 1—Terminals
- PU Type 2—Controllers
- PU Type 3—not defined
- PU Type 4—Front-end Processors
- PU Type 5—Hosts

### Logical Units

Logical units (LUs) are software-based access points through which users can communicate with the rest of an SNA network. Each LU in the following list represents functions that software programs provide in an SNA network:

- LU0—generic function
- LU1—printer support
- LU2—3270 screen management
- LU3—3270 printer management
- LU6.2—program-to-program communication (used for APPC/APPN)

### System Service Control Points

Simply put, system service control points (SSCP) is the software that provides the overall management services for a particular domain of the SNA network. Obviously, without this service, chaos would abound on an SNA network. System service control points are found on Type 5 nodes or host processors.

### Network Addressable Units

Each node contains one or more network addressable units (NAUs) and path control network components. PUs, LUs, and SSCP are all network addressable units. In addition to the address, these units are given a network name (or alias name) that can be easier to remember than a set of hexadecimal numbers. For example, an IBM 3820 printer has a unique network address and name. The name might be representative of the city and floor where the printer is located (i.e., HOU22). The address will be its hexadecimal address.

**SNA (*continued*)****Domain**

A domain is a collection of SNA nodes in a network that are managed by an SSCP. This SSCP management is a soup-to-nuts management service including devices, physical wiring, software, and microcode. Domains typically contain several subareas in a medium- to large-size SNA network.

**Subareas**

A subarea can be considered a branch of the domain in a hierarchical topology. Each subarea within a single domain communicates with the same host.

**Multiple Domain Networks**

When more than one host (Type 5 device) is present, multiple domains exist. Devices in domain A, for example, can communicate with devices in domain B. A host in domain A and a host in domain B set up communications that allow for devices in the two domains to communicate as well.

**The Pictorial View**

Sometimes you can describe a concept until it's dead and your readers may still not understand what you're trying to convey. We believe the old adage about a picture being worth a thousand words. Figure 2 depicts all the items mentioned in the preceding sections in a simplistic setup that should help you understand their functions and relationships.

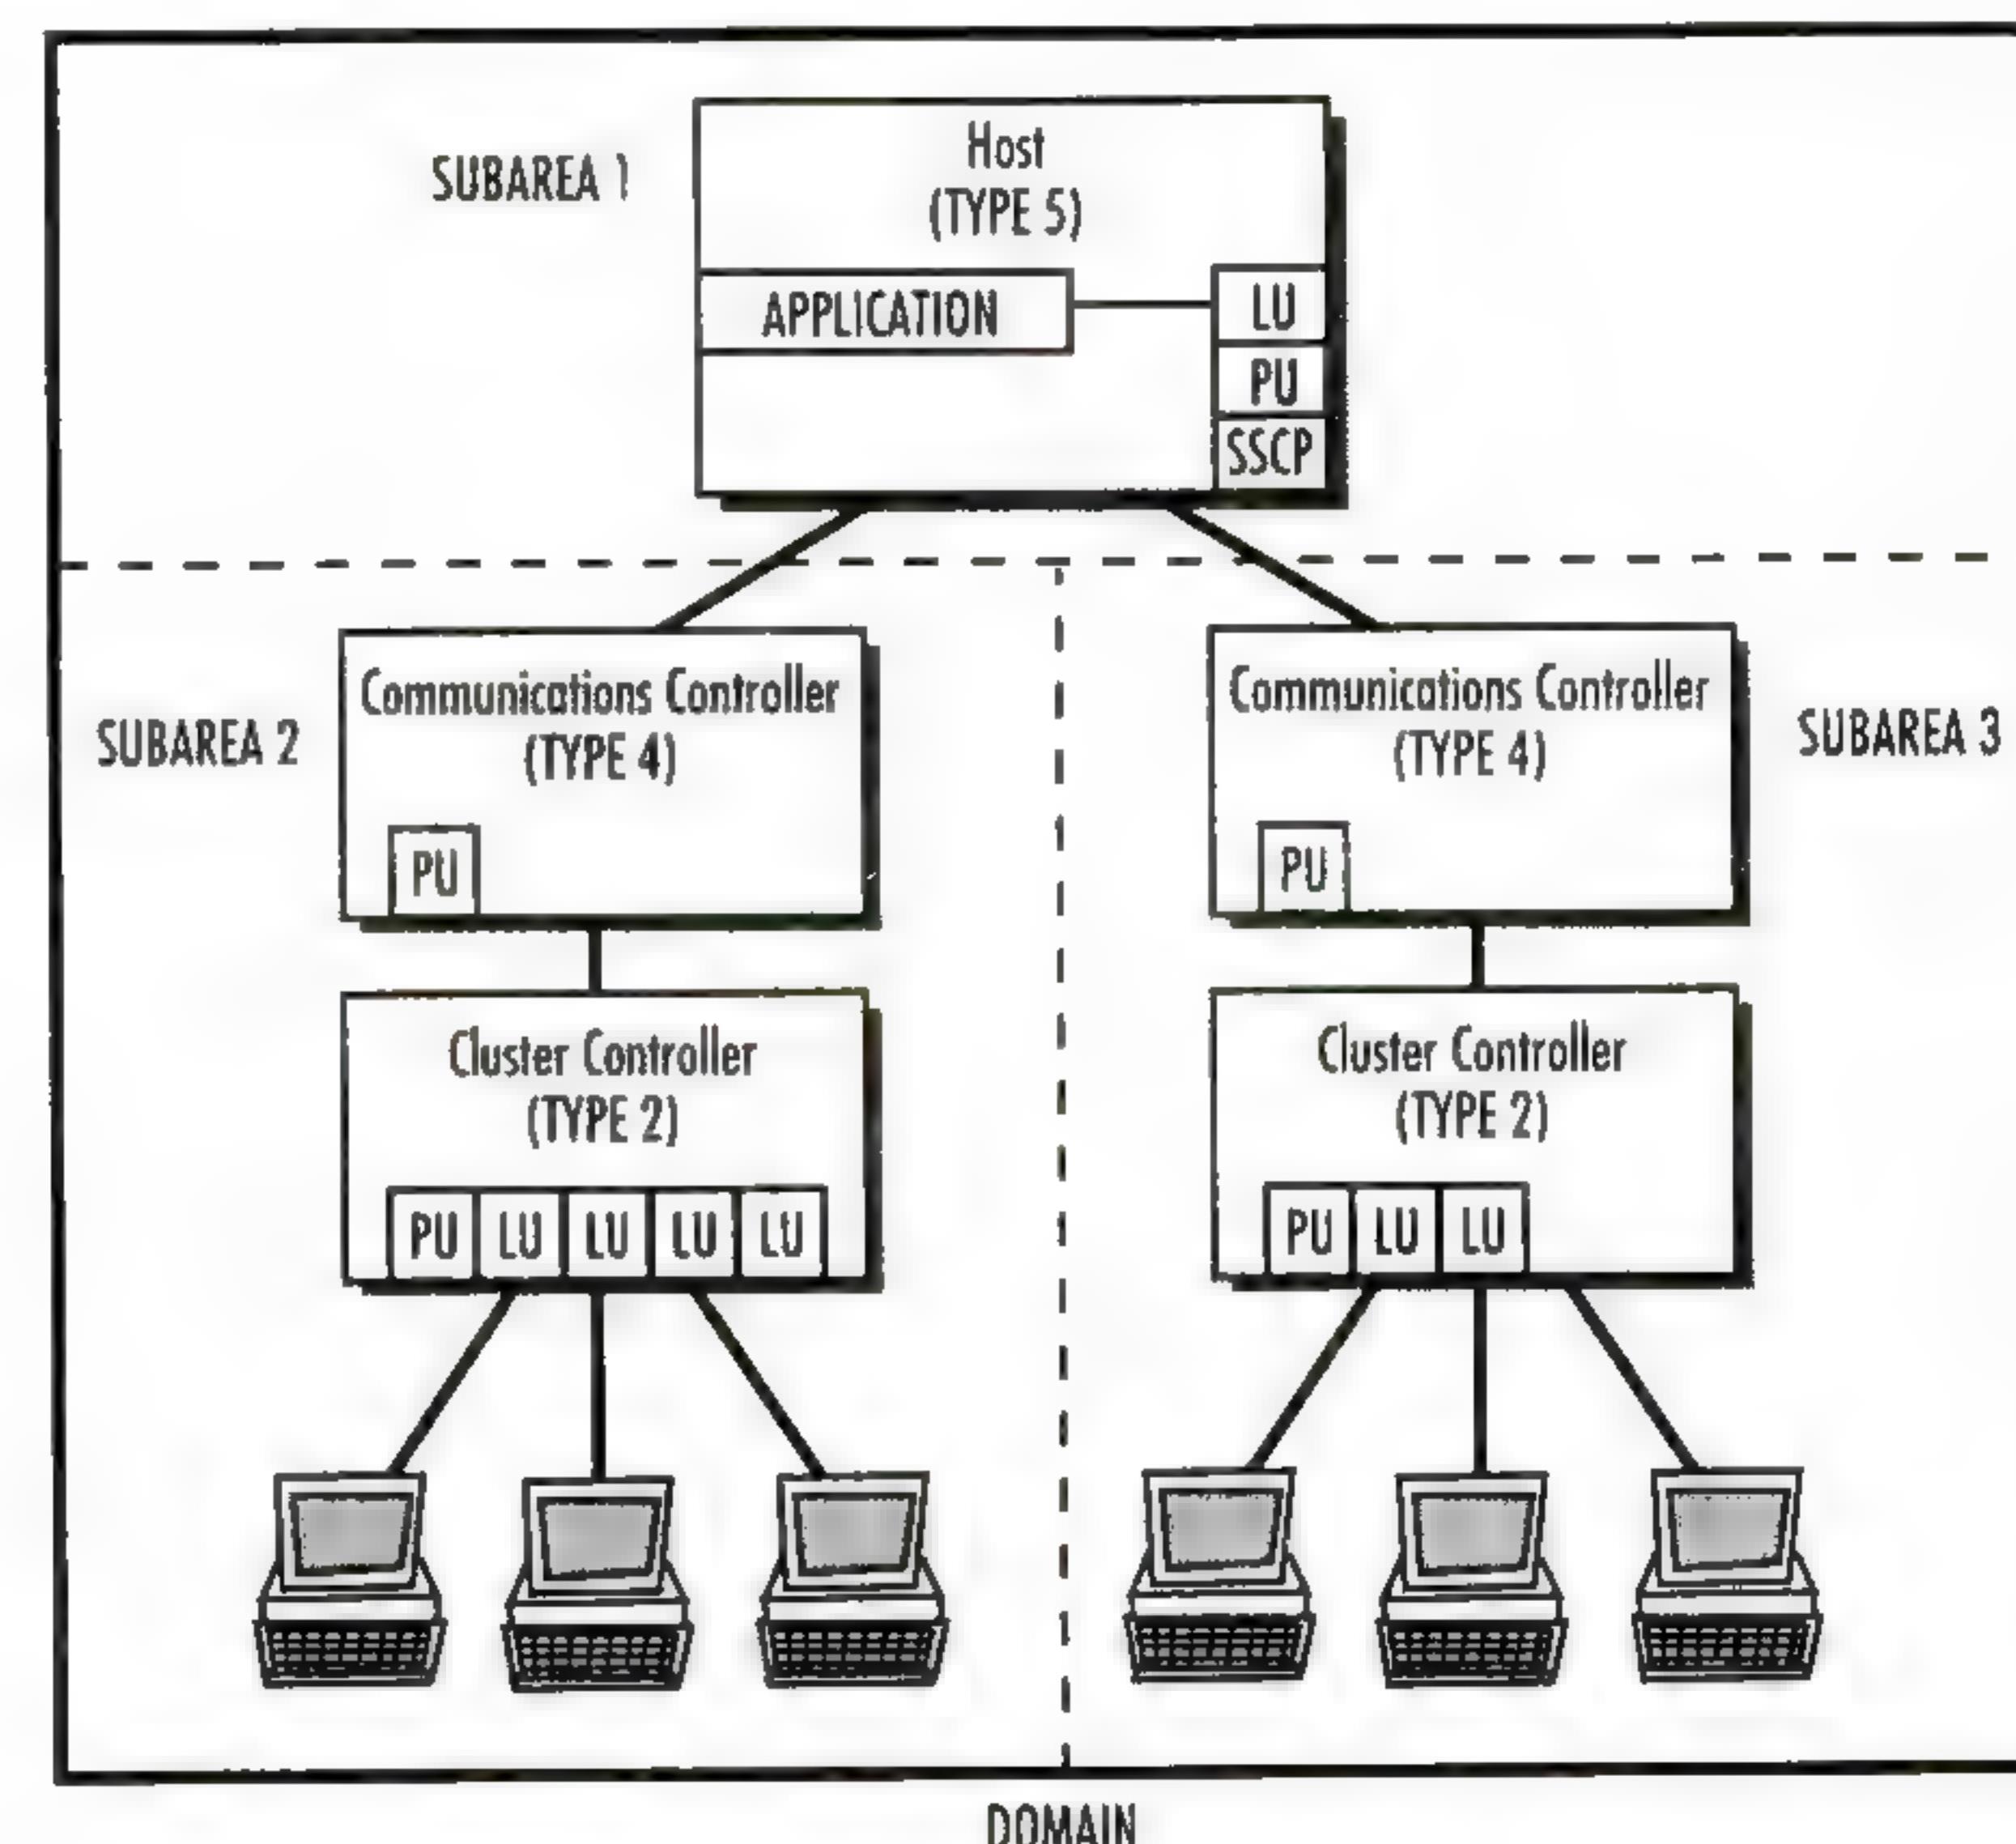


Figure 2: A basic SNA network

**SNA vs. the OSI Model**

Figure 3 shows a comparison of the SNA architecture to the OSI model. The SNA model only has five layers compared to the seven layers of OSI. Even though layers 1 and 7 of the OSI model are important even in the SNA world, they are defined outside the SNA model. Nevertheless, we've placed them here as a reference.

Layer	OSI Model	SNA Model
7	Application	Application
6	Presentation	Function Management
5	Session	Data Flow Control
4	Transport	Transmission Control
3	Network	Path Control
2	Data Link	Data Link Control
1	Physical	Physical Control

Figure 3: OSI Model versus SNA Model

## A Layered Approach

Each layer in the SNA model is assigned functions to perform and services to provide. We'll go through each layer individually, but first we want you to look at the layers as a whole, and see how each layer communicates with the others.

Figure 4 shows two basic nodes. Node 1 would like to communicate with node 2. The user at node 1 passes the information on to the function management layer. That layer may add information to the message, as determined by its function, and then will pass the data down to the next layer—data flow control. Each layer continues to add information and passes its data to the next layer below until the physical control layer is reached. When the data reach this layer, they are placed onto the wiring and sent across to node 2.

Node 2 reverses the process by sending the data up through the layers until the data reach the user at node 2. As the data are passed up through node 2's layers, each layer strips off any information that its function may have added in node 1. By the time the data are received by the user at Node 2, all extra information is stripped off, leaving only the data. This description may make the layer information seem extraneous, but it's vital to the accurate packaging and delivery of data across the network from one node to another.

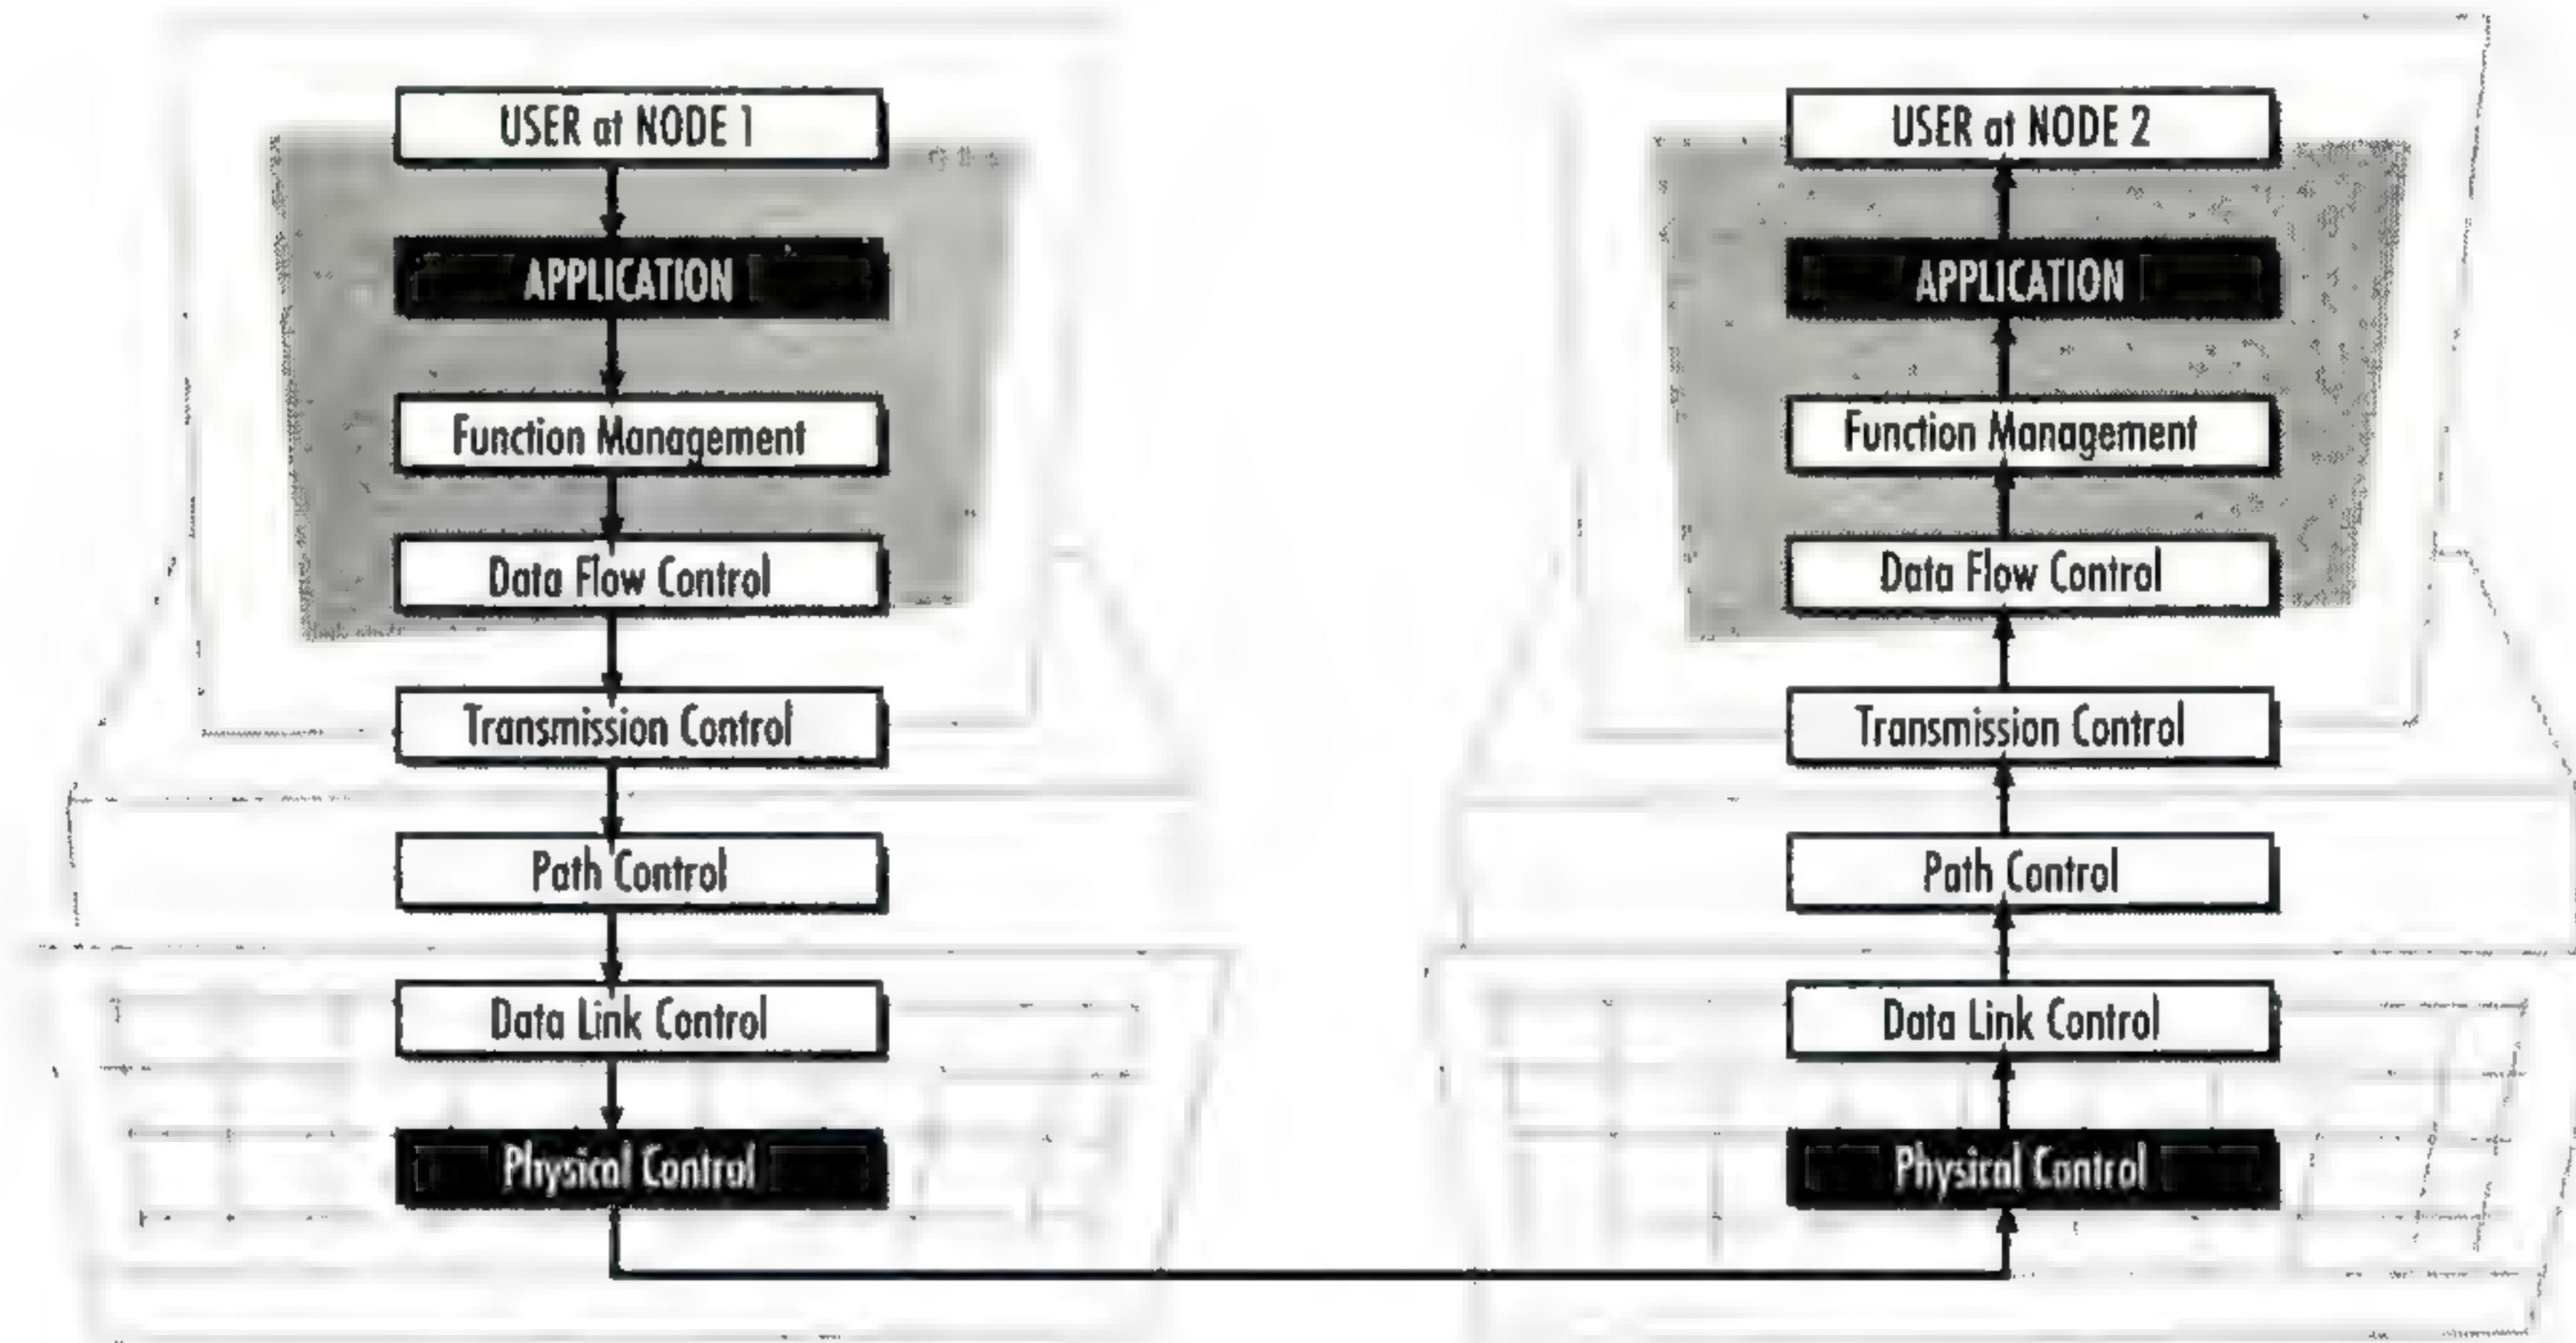


Figure 4: Node-to-node communication

### Physical Control Layer

The physical control layer is defined outside of the SNA architecture; like the physical layer in the OSI reference model, it is concerned with sending bits over the wire. This layer doesn't concern itself with interpreting bits or trying to decipher how many bits are grouped together. It simply knows that it has some bits that need to get onto the wire but doesn't check to see where the information needs to go.

### Data Link Control Layer

The data link control layer operates one level above the physical control layer. This layer concerns itself with transmitting the data from one node to the next with the routing information provided by the path control layer. The data link control layer is also responsible for detecting and recovering from transmission errors that may occur on the physical link.

### Path Control Layer

This layer is responsible for providing routing information for the data. This layer defines the end-to-end path that a message must take to the data link control layer. The path might include the various nodes that a message must visit to get from its sending to its destination address.

*continued on next page*

**Transmission Control Layer**

The transmission control layer maintains session status and tracks packet sequencing, data integrity, and data pacing. Once a session is initiated, this layer keeps an eye on the session's progress. It also looks at whether the data are received in the same sequence they were sent. If the sending node encrypts the data, the transmission control layer decrypts the data on the receiving end.

**Data Flow Layer**

The data flow layer is charged with managing the overall flow of data during a session.

**Function Management Layer**

The function management layer is too complex for one layer, so it is functionally divided into two sublayers:

- Function management data services and
- The NAU services manager

The function management data services sublayer is the manager of the interface provided to the user as well as the presentation of that interface. A completely separate function of this sublayer is to manage the overall network and its associated sessions.

The NAU services manager layer is the nebulous layer. It is usually described only as a layer that provides services to layers beneath it. We're not exactly sure what IBM does here, and neither are any of the references we've checked.

The application layer is not part of the SNA model, but is defined outside of the SNA architecture. This layer represents the users of the network and the network software.

Today it is common to find SNA users connected over a LAN through a gateway server to the mainframe world. Novell has been the leader in this technology with its NetWare for SAA product. Microsoft NT now offers a remote office gateway product that matches Novell's product (and, some would say, exceeds it in ease of use and functionality). Regardless, a gateway generally has options to connect to the SNA network using SDLC or the channel attach method. Because this forces network administrators to support two backbones, this is a desirable interim step.

The real challenge involved in internetworking the enterprise into one homogeneous network is to provide network services with only one backbone technology, not several. Network management and security are critical issues in this challenge. Some router vendors include SNMP-to-NetView mapping modules to overcome this problem by interlinking TCP/IP-based SNMP with IBM's NetView network management environment. A tightly monitored SNA network can look scary when viewed through a TCP/IP environment and vice versa. Vendors are now attempting to address security issues to overcome this problem.

**LAN Traffic vs. SNA Traffic**

So, what's all the hubbub about connecting SNA networks over an enterprise backbone? If the enterprise backbone carries LAN traffic, SNA nodes could run into some congestion problems. Because LAN traffic can be classified as "bursty" traffic, meaning that there are times when the pipe is quiet and times when the entire bandwidth of the pipe might be utilized, this can pose a serious problem for time-sensitive SNA communications.

**SNA (*continued*)**

For example, if an FEP shares the same backbone with a file server on a LAN, the FEP may not be able to communicate, or poll its controllers, if a user is transferring a huge graphic file across the backbone. Until recently, this problem could be solved only by separating the two types of networks, using separate and distinct backbones. This solution is not only complex and but requires organizations to maintain trained support personnel for each type of technology. Clearly, a better common solution was needed.

### **Vendors and Groups Lead the Way**

Because IBM has been slow to release standards and strategies to help SNA networks integrate with LANs, vendors have been slow to provide products. Cabletron, Hewlett-Packard, 3Com Corp., and CrossCom, Inc. are router vendors that have announced support for something called the *DLSw* specification (data link switching, a backbone technology that mixes support for guaranteed bandwidth required for SNA or other time-sensitive traffic with bandwidth on demand suitable for LAN traffic). These vendors were involved in an interoperability test where IBM rigorously tested the DLSw specification with various products from the vendors.

In October 1994, an *APPN Implementors Workshop* (AIW) was held. The AIW is involved in finalizing its high-performance routing (HPR) standard. HPR is a high-speed standard for multiprotocol backbones in enterprise networks. AIW is also involved in the completion of the DLSw standard mentioned earlier. It's still too early to tell what contributions this group will be able to make, if any, but its very existence augurs well for IBM LAN-to-host connectivity issues.

### **Data Link Switching**

IBM originated this informal specification and implemented it in the 6611 router. The DLSw specification defines a way to send SNA traffic over TCP/IP-based backbones. IBM then submitted DLSw to the Internet Engineering Task Force as a proposed standard, RFC 1434, in 1993. At present, this standard is nearing completion and will probably be an official standard by the time you read this.

DLSw allows routers to protect the time-sensitive nature and data integrity of SNA networks and hence should be adopted by most major router vendors.

### **Frame Relay**

Frame Relay is well suited to respond to the challenge of integrating SNA and LAN-based traffic. Frame Relay requires less error checking, which can slow down other transmission methods, and it works well in a "bursty" environment. Because SNA provides its own error checking, it can combine elegantly with Frame Relay. Because Frame Relay also handles "bursty" traffic well, it appears able to serve both sets of networking concerns more or less equally.

In addition, Frame Relay has gained popularity as a means for SNA internetworking because it is less expensive than leased lines and can achieve higher speeds. It is a fast-growing industry that is expected to replace a lot of the leased lines and X.25 network backbones.

### **ATM**

IBM has chosen Asynchronous Transfer Mode (ATM) as its future internetworking direction. ATM technology is well suited for data, voice, and video because it uses small, fixed-length packets called *cells*. ATM's most attractive feature for the SNA community is that it will include "bandwidth on demand" services. It allows organizations to pay for the transmission bandwidth they actually consume, instead of paying a fixed monthly cost for a pre-specified level of bandwidth on a leased line. All of these should help prolong the usefulness of SNA.

**SNA (*continued*)****TCP/IP**

TCP/IP appears to be the predominant protocol for internetworking enterprises. A recent push in this direction is IBM's draft proposal of DLSw, which is designed to route NetBIOS and SNA traffic over TCP/IP-based networks. Many of the router companies are already implementing the DLSw specification in their router technologies.

**Summary**

SNA continues to provide networking connections and technologies for many organizations, especially those who have been involved with networks the longest. Because so many of these organizations have substantial investments in SNA, the protocol appears destined to stick around for the foreseeable future. As these organizations grapple with the problems in internetworking SNA- and LAN-based networks, they have pushed IBM and the networking industry toward a variety of interesting and workable solutions. We've only scratched the surface of SNA in this article, though, so we invite you to consult the information in the following section, if you're in need of further details.

**References**

- [1] Martin, James, *SNA: IBM's Networking Solution*, Prentice-Hall, 1987. [This book thoroughly describes the SNA architecture and is geared toward the technical guru; novices beware.]
- [2] Cerutti, Daniel, & Donna Pierson, *Distributed Computing Environments*, McGraw-Hill. [A good summary of the SNA architecture can be found within a few pages of this book.]
- [3] Pedersen, Elinor, "The APPN Connectathon," *Midrange Systems*, Volume 7, No. 15 (August 12, 1994). [This is a great overall article on APPN.]
- [4] Krivda, Cheryl D., "Where the Action Is," *Internetwerk*, Volume 5, No. 11 (November 1994). [This is an excellent article to read if you are attempting SNA internetworking. The article includes address and phone numbers of vendors that provide SNA internetworking products.]
- [5] Clark, Wayne, "SNA Internetworking," *Connexions*, Volume 6, No. 3, March 1992.
- [6] Joyce, Steven T. and Walker II, John Q., "Advanced Peer-to-Peer Networking (APPN): An Overview," *Connexions*, Volume 6, No. 10, October 1992.
- [7] Clark, Wayne, "Accommodating SNA Peer-to-Peer Networking in a Multiprotocol Environment," *Connexions*, Volume 7, No. 3, March 1993.
- [8] "APPN Architecture and Product Implementations Tutorial," IBM Publication No. GG24-3669-01, June, 1992.
- [9] R. Dixon, D. Kushi, "Data Link Switching: Switch-to-Switch Protocol," RFC 1434, March 1993.

[This article is based on material in *The PC Networking Handbook*, by Ed Tittel, ISBN 0-12-691398-6, published by Academic Press Professional. Used with permission. —Ed.]

**ED TITTEL** is a full-time freelance writer and networking consultant, and a member of the NetWorld+Interop Program Committee. Ed is the author of ten books on computer topics, many of them network-related. In addition to *NetWare for Dummies*, now in its 2nd edition, Ed is the author of the forthcoming IDG books *HTML for Dummies* and *Foundations of WWW and HTML Programming*, as well as books on the Internet, network design, and e-mail for Academic Press. In his spare time, Ed likes to walk his hefty but personable Labrador, Dusty, and to share the blessings of domesticity with his wife, Suzy, and stepchildren, Austin and Chelsea. His e-mail address is: [etittel@zilker.net](mailto:etittel@zilker.net)

## Call for Papers

The *1996 International Conference on Network Protocols* (ICNP-96) will be held October 29–November 1, 1996 at the Hyatt Regency Hotel, Columbus, Ohio, USA. ICNP-96 is sponsored by the IEEE Computer Society Technical Committee on Distributed Processing (TCDP), in cooperation with the Information Processing Society of Japan (IPSJ).

**Topics** Original technical papers addressing the following topics of interest are solicited for presentation at the conference and publication in the conference proceedings:

- Network Architecture
- Routing Protocols
- Real-Time Protocols
- Protocol Conversion
- Communication Software
- Network Service
- Internetworking
- Multimedia Systems
- Protocol Design Methodology
- Protocol Verification/Validation
- Distributed Cooperative Systems
- Protocol Testing and Debugging
- Flow & Congestion Control
- Wireless & Mobile Networks
- High-Speed Networks
- Network Management
- Broadcast Systems
- Network Security
- Group Communication
- ATM Protocols
- Transport Protocols
- Switching Protocols
- Protocol Implementation
- Intelligent Communication

**Submissions** Authors are requested to submit six copies (in English) of their double-space typed manuscript (maximum of 25 pages) with an abstract to the program chair by April 1, 1996. All accepted papers will be published in the conference proceedings; some will be forwarded to the journal *IEEE/ACM Transactions on Networking* for publication consideration. Submit papers to:

Prof. Hasan Ural, Program Chair  
 Department of Computer Science  
 University of Ottawa  
 MacDonald Hall  
 150 Louis Pasteur  
 Ottawa  
 Ontario K1N 6N  
 CANADA  
 Tel: +1 (613) 562-5800 ext. 6684  
 Fax: +1 (613) 562-5185  
 E-mail: [ural@csi.uottawa.ca](mailto:ural@csi.uottawa.ca)

**Tutorials** The day preceding the conference (October 29) will be devoted to tutorials to provide the background for the conference. Send tutorial proposals by April 1, 1996 to:

Prof. Ten H. (Steve) Lai  
 Dept. of Computer and Information Science  
 The Ohio State University  
 2015 Neil Avenue  
 Columbus, Ohio 43210-1277  
 U.S.A.  
 Tel: +1 (614) 292-2146  
 Fax: +1 (614) 292-2911  
 E-mail: [lai@cis.ohio-state.edu](mailto:lai@cis.ohio-state.edu)

<b>Important dates</b>	Papers and proposals due:	April 1, 1996
	Acceptance letters sent:	July 1, 1996
	Camera-ready copies due:	August 15, 1996

## Call For Papers

The *IEEE Global Internet 1996* conference takes place during *Globe-com '96* and will be held at the Queen Elizabeth II Conference Centre in London, November 20 and 21, 1996.

### Goals

This mini-conference will be jointly organized by the two IEEE Communications Society Technical Committees on Internet and Computer Communications. It will provide an open forum for the communications and computer networking communities to review the state-of-the-art technologies and applications of the evolving Global Internet. It will also provide an opportunity to highlight solutions to pressing issues, establish a vision for the future, and challenge the participants to press forward in their research and engineering efforts to meet business and industry needs for global internetworking.

The mini-conference will include keynote speeches, tutorials and panel discussions by leaders in these areas, invited papers by recognized experts, and contributed papers by active researchers in the field. The conference will put special emphasis on hands-on experiences with actual implementations and widespread applications.

### Topics

Submissions should be on key topics, including the following:

- *General topics:*
  - Evolution of the Internet and WWW: past, present, and future
  - Legal and regulatory issues (e.g., censorship)
  - Privacy, security and billing
- *WWW technology and applications:*
  - Protocol evolution and extensions
  - Tools and browsers
  - Authoring environments
  - Retrieval and resource discovery
  - Information representation, modeling and filtering
  - Consistency, integrity and security
  - User and application interfaces
  - Nomadic software
  - Virtual reality
  - Integration of real time data
  - Design techniques for Web applications
  - Kiosk systems
  - Computer based training, teaching and CSCW
  - WWW applications for corporate “intranets”
- *Internetworking technologies and applications:*
  - Routing, addressing, naming, and large scale multicast
  - ATM, Frame Relay and SMDS as part of the Internet
  - Performance and reliability
  - Nomadic computing and communications
  - Multimedia services to the desktop
  - Interactive multimedia video, sound etc. on commercial nets
  - Internet Telephony
  - Distributed Simulation
  - Internet management experiences and solutions
  - Enterprise networking architectures and applications
  - Broadband consumer/home access to the Internet
  - Virtual Corporate Networks
  - Security, billing, and privacy for electronic commerce

**How to submit papers**

Electronic submissions are preferred. If your browser supports file upload (Netscape 2.0 does), you can use the form at:

<http://gaia.cs.umass.edu:80/tccc/internet96/>

If you cannot use upload, please e-mail *PostScript* to either:

Jon Crowcroft ([jon@cs.ucl.ac.uk](mailto:jon@cs.ucl.ac.uk))

or

Henning Schulzrinne ([schulzrinne@fokus.gmd.de](mailto:schulzrinne@fokus.gmd.de))

If you cannot submit your paper in electronic form, mail it to:

Professor J. Crowcroft  
Department of Computer Science  
UCL  
Gower Street  
London WC1E 6BT  
ENGLAND

*Submissions by fax are not accepted.*

**Important dates**

May 15th, 1996:	Deadline for contributed paper submission
July 15th, 1996:	Notification of paper acceptance to authors
September 1, 1996:	Revised manuscript due

**Panel discussions**

- Internet Exporting, Trading, Marketing
- Java
- The Global Internet: Technical Progress in Worldwide Inter-networking

**Tutorials**

- Interactive Multimedia over the Internet (Mark Handley/Ian Wakeman)
- Java (John Daigle)
- Internet Security (Ran Atkinson)
- IP/ATM (Dave Ginsberg)

**Conference Committee**

Brian Carpenter, General Chair  
Jon Crowcroft, Technical Chair  
Henning Schulzrinne, Vice Technical Chair  
Roch Guerin, Vice Technical Chair

**Technical Committee**

Grenville Armitage, Fred Baker, Bob Braden, Andrew T. Campbell, Brian Carpenter, Nim K. Cheung, Jon Crowcroft, Laura Cunningham, John N. Daigle, Sally Floyd, Raj Jain, Srinivasan Keshav, Jim Kurose, Craig Partridge, Henning Schulzrinne, Jonathan M. Smith, Joe Touch, John Wroclawski, Yechiam Yemini, Lixia Zhang.

## Preliminary Call for Participation

The *Fourth ACM International Multimedia Conference and Exhibition* (ACM Multimedia '96) will be held November 18–22, 1996 at the Hynes Convention Center in Boston, MA, USA. The events is sponsored by the ACM SIG Multimedia, SIGCOMM, SIGLINK, SIGMIS and SIGGRAPH, in cooperation with SIGBIT, SIGIR, SIGCHI, and SIGOIS.

Multimedia technology can substantially improve the communication between information providers and consumers. It contributes to the general accessibility of information, through new interactive media as well as through new forms of production, delivery and perception of existing media. ACM Multimedia '96 will provide an international forum for papers, panels, videos, demonstrations, courses, workshops, and exhibits focusing on all aspects of this multi-disciplinary field: from underlying technologies to applications and issues, and from theory to practice. We invite your participation.

### Topics

Topics include, but are not limited to: applications in art, education, entertainment, government, medicine, etc.; collaboration environments; databases; digital libraries; distributed systems; documents and authoring; hardware and architectures; image, video and audio compression techniques; information retrieval; interactive television; media integration and synchronization; networking and communication; operating system extensions; programming paradigms and environments; standards and legal issues; storage and I/O architectures; tools; user interfaces; and virtual reality.

ACM Multimedia '96 will be co-located with SPIE's Symposium on Voice, Video and Data Communications, and Broadband Communications Expo. It will overlap with CSCW, to be held in nearby Cambridge.

### Papers

Technical papers on completed or in-progress research, innovative applications, or experience with multimedia systems are solicited. Submissions must use a minimum of 10-point typeface, and be up to 12 pages (preferably double sided), including figures, tables, and references. Where applicable, prototype demonstrations or videotape presentations are encouraged to supplement the papers. Papers must be accompanied by an electronic cover sheet (see submission instructions below). Submit complete papers to: Wendy Hall, Program co-Chair ([W.Hall@ecs.soton.ac.uk](mailto:W.Hall@ecs.soton.ac.uk)).

Outstanding papers on different areas of multimedia will be given awards. Papers with a student as the primary author will enter a student paper award competition. A cover letter must identify the paper as a candidate for the student paper competition. Selected papers will be forwarded to ACM/Springer-Verlag *Multimedia Systems*, *Communications of the ACM*, IEEE/ACM *Transactions on Networking*, or ACM *Transactions on Information Systems*.

### Multimedia and Art

Submissions by artists presenting innovative work in the field are encouraged. A specific selection process and a special Multimedia and Art session will take place. Submissions by artists should include a paper presentation, a single VHS NTSC video and demonstration requirements when applicable, and a biography. Submit to: Art chair. (contact to be announced, see Web page for updates).

<b>Panels</b>	Panels are solicited that examine innovative, controversial, or otherwise provocative issues of interest. Proposals should be limited to 2 pages, plus a biography of at most one paragraph for each participant. Submit proposals to: Bob Allen, Panels Chair ( <a href="mailto:rba@bellcore.com">rba@bellcore.com</a> ).
<b>Demonstrations</b>	We solicit demonstrations of working systems in technical and artistic categories. Submissions (at most 2 pages) should include a description of the exhibit, demo requirements, a biography, and a single VHS NTSC video. Submit demonstrations to: Arding Hsu, Demonstrations Chair ( <a href="mailto:ahsu@scr.siemens.com">ahsu@scr.siemens.com</a> ).
<b>Courses</b>	There will be a series of 1/2-day tutorial courses, focused on issues relevant to researchers and/or practitioners of multimedia technology. Proposals (at most 5 pages) should include a description of the subject matter and brief biographical sketches of the instructors. Evaluation of proposals will be based on expertise and experience of instructors, relevance of subject matter, and the use of multimedia technology in the presentation. Submit tutorial proposals to: Rajiv Mehrotra, Tutorials Chair ( <a href="mailto:rajiv@mayura.cs.ums1.edu">rajiv@mayura.cs.ums1.edu</a> ).
<b>Workshops</b>	Workshops preceding the conference will allow participants to exchange ideas on a topic. Workshop results and issues will be integrated into the main body of the conference. Submit workshop proposals to: Wayne Wolf, Workshops Chair ( <a href="mailto:wolf@ee.princeton.edu">wolf@ee.princeton.edu</a> ).
<b>Exhibits</b>	Exhibits for ACM Multimedia '96 will be combined with those for SPIE's Symposium on Voice, Video, and Data Communications, offering vendors and publishers a unique opportunity to exhibit and demonstrate multimedia products. For more information, contact <a href="mailto:exhibits@spie.org">exhibits@spie.org</a> .
<b>Submissions</b>	Authors should consult the World-Wide Web at URL:  <a href="http://www.acm.org/sigmm/MM96/">http://www.acm.org/sigmm/MM96/</a> ...for more detailed submission guidelines and up-to-date information about ACM Multimedia '96. Authors of accepted submissions will be required to submit both a camera-ready copy of the manuscript for the printed proceedings and an electronic copy for the CD ROM proceedings. Authors must assign copyright to ACM as a condition of publishing their work in the proceedings. An author who embeds an object, such as an art image, copyrighted by a third party is expected to obtain that party's permission to include the object with the understanding that the entire work may be distributed as a unity to ACM members and others.
<b>Important Dates</b>	All Submissions (6 copies for papers) due: April 24, 1996 Notification of acceptance: July 15, 1996 Final submissions due: August 26, 1996
<b>More Information</b>	<a href="http://www.acm.org/sigmm/MM96/">http://www.acm.org/sigmm/MM96/</a> or <a href="http://www.uni-mannheim.de/acm96/">http://www.uni-mannheim.de/acm96/</a>

## Book Review

*Active Java: Object Oriented Programming for the World-Wide Web*, by Adam Freeman and Darrel Ince (OU), Addison-Wesley (Longman), Jan 1996, ISBN 0-201-40370-6, 235 pages

### Introduction

This is one of a small number of books on the new programming language *Java*, from Sun Microsystems (and now implemented in the Netscape Navigator 2 beta web browser, as well as by a number of other compiler/language support companies, notably Borland).

There are several other books in the market now (and no doubt there will soon be hundreds), although this reviewer has only found one that is actually worth getting, which is *Teach yourself Java in 21 Days* from SAMS, which is encyclopedic, and quite expensive (and features a number of typographical errors, although it is largely very accurate, and far more complete than other books).

This new book, however, has a somewhat different goal, which is more suited to introducing people to a programming language without necessarily assuming that they have learned programming before (e.g., most other books assume **C/C++** familiarity, of Visual Basic/**VC++**, or even *Perl/Python/Awk/Tcl/Tk* etc., etc.).

### Organization

The book comprises 2 parts, each in 6 chapters. The first part is about the base language concepts, ranging from the goals of Java and interactive Internet applications, through Objects, Classes, basic data types and control structures, Java classes and libraries. The second part covers some specific class libraries, such as the *Abstract Window Toolkit* (AWT), the networking class library, the *Java Development Kit* (JDK), and building an Applet (a class that is run from a web page) and building an application; the last chapter looks at some Java internals, such as the portable byte code (like *Pascal* (e.g., UCSD *p-code*), just-in-time compiling, the virtual machine, Java security, Applets and applications and Java futures.

The style of the book is readable, and informal, and filled with many small examples, both code fragments, and small complete programs. A lot more material is promised on the follow-up web site, which, given the authors' employer is the UK Open University, will certainly appear—the publisher is making some of the material directly available:

<http://www.aw.com/cseng/authors/freeman.a/activejava/actjava.html>

### Critique

The book is a tad early, since a number of things about Java, particularly in the area of scoping, inheritance (multiple inheritance of interfaces, but not of classes), security (byte code machine safety has been called into question by several papers—e.g., by Dan Wallach from Princeton University, Computer Science Department) are unsettled. However, most users (this reviewer included) have had extremely positive experiences using JDK recently, and believe that it will certainly go a long way, possibly even supplanting **C/C++** eventually as a safer, and easier to learn language.

### Audience for this book

The curious engineer or student could use this book to catch up quickly with Java, but also it might act as a base book from which to teach, say, an undergraduate second programming course on Java—certainly, right now, there is no better book (though for postgraduates, or the development engineer, the SAMS book is essential).

As the authors point out, most of the technical details of JDK and the language are freely available on the net, although they are hard to gather together (and in these days of congested super-lowway, often too slow to access unless you are lucky enough to have a good cache or mirror site near you). However, nothing quite beats the printed page for focussed learning.

**Problems**

The book doesn't have enough expanded examples—it would be nice to have a real OU teaching example (a computer aided learning package—on anything, e.g., cookery, graphics, juggling, you name it!) The book doesn't have a references/citations/further reading section—a list of URLs at least would be nice! The book doesn't list the known security problems (with Netscape, Hotjava, Javap, the byte code language, etc., etc.). The book could also do with a syntax summary (appendix).

**Summary**

7 out of 10—a useful text, in a sea of useless bandwagons with only one lifeboat (the SAMS book) so far. The second edition will be important!

—*Jon Crowcroft, University College London*  
J.Crowcroft@cs.ucl.ac.uk

---

**Write to *ConneXions*!**

We'd love to hear your comments, suggestions and questions about anything you read in *ConneXions*. Our editorial address is given below. Use it for letters to the Editor, requests for the index of back issues, questions about particular articles etc.:

*ConneXions—The Interoperability Report*  
303 Vintage Park Drive  
Suite 201  
Foster City  
California 94404-1138  
USA  
Phone: +1 415-578-6900 or 1-800-INTEROP (Toll-free in the USA)  
Fax: +1 415-525-0194  
E-mail: [connexions@interop.com](mailto:connexions@interop.com)  
URL: <http://www.interop.com>

**Subscription information**

For questions about your subscription please call our customer service hotline: 1-800-575-5717 or +1 610-892-1959 outside the USA. This is the number for our subscription agency, Seybold Publications. Their fax number is +1 610-565-1858. The mailing address for subscription payments is: P.O. Box 976, Media, PA 19063-0976.

---

This publication is distributed on an "as is" basis, without warranty. Neither the publisher nor any contributor shall have any liability to any person or entity with respect to any liability, loss, or damage caused or alleged to be caused, directly or indirectly, by the information contained in *ConneXions—The Interoperability Report*®

---

**CONNE~~X~~IONS**

303 Vintage Park Drive  
 Suite 201  
 Foster City, CA 94404-1138  
 Phone: 415-578-6900  
 FAX: 415-525-0194

FIRST CLASS MAIL  
 U.S. POSTAGE  
 PAID  
 SAN JOSE, CA  
 PERMIT NO. 1

ADDRESS CORRECTION  
 REQUESTED

**CONNE~~X~~IONS**

**EDITOR and PUBLISHER** Ole J. Jacobsen

**EDITORIAL ADVISORY BOARD** Dr. Vinton G. Cerf  
 Senior Vice President, MCI Telecommunications  
 President, The Internet Society (1992 – 1995)

A. Lyman Chapin, Chief Network Architect,  
 BBN Communications

Dr. David D. Clark, Senior Research Scientist,  
 Massachusetts Institute of Technology

Dr. David L. Mills, Professor,  
 University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,  
 University of Southern California, Information Sciences Institute



Printed on recycled paper

**Subscribe to CONNE~~X~~IONS**

**U.S./Canada**  \$195. for 12 issues/year

**All other countries**

\$245. for 12 issues/year

Name \_\_\_\_\_ Title \_\_\_\_\_

Company \_\_\_\_\_ E-mail \_\_\_\_\_

Address \_\_\_\_\_

City \_\_\_\_\_ State \_\_\_\_\_ Zip \_\_\_\_\_

Country \_\_\_\_\_ Telephone (      ) \_\_\_\_\_

Fax (      ) \_\_\_\_\_

Check enclosed (in U.S. dollars made payable to CONNE~~X~~IONS).

Visa  MasterCard  American Express  Diners Club Card# \_\_\_\_\_ Exp.Date \_\_\_\_\_

Signature \_\_\_\_\_

**Please return this application with payment to:**

**CONNE~~X~~IONS**

Back issues available upon request \$15./each  
 Volume discounts available upon request

303 Vintage Park Drive, Suite 201  
 Foster City, CA 94404-1138  
 415-578-6900 FAX: 415-525-0194  
**connexions@interop.com**